National Center for Education Statistics. (2003). A profile of participation in distance education: 1999–2000 (NCES 2003-154). Washington, DC: US Department of Education, Office of Educational Research and Improvement. Available: http://nces.ed.gov/pubs2003/2003154.pdf.

NCS secures rights to iPaper electronic technology in testing and education market. (2000, July 11). Minneapolis, MN: National Computer Systems (NCS). Available: www.ncs.com/ncscorp/top/news/000711.htm.

Negroponte, N. (1995). *Being digital*. New York: Vintage.

Odendahl, N. (1999, April). *Online delivery and scoring of constructed-response assessments*. Paper presented at the annual meeting of the American Educational Research Association, Montreal.

Olsen, F. (2000, February 18). A UCLA professor and net pioneer paves the way for the next big thing. *The Chronicle of Higher Education, 46*.

Olson, L. (2003). Legal twists, digital turns: Computerized testing feels the impact of "No Child Left Behind." *Education Week, 12*(35), 11–14, 16.

Pellegrino, J. W., Jones, L. R., & Mitchell, K. J. (1999). *Grading the nation's report card*. Washington, DC: National Academy Press.

Perry, J. (2000). Digital tests spark controversy: Critics say revamped exams limit the options to challenge a score. *Online US News* [On-line]. Available: www.usnews.com/usnews/edu/beyond/grad/gbgre.htm.

Russell, M., & Plati, T. (2001). Effects of computer versus paper administration of a state-mandated writing assessment. *Teachers College Record*. Available: www.tcrecord.org/Content.asp?ContentID=10709.

Schmidt, P. (2000, January 21). Judge sees no bias in Texas test for high-school graduation. *Chronicle of Higher Education*, p. A27.

Schulze, K. G., Shelby, R. N., Treacy, D. J., & Wintersgill, M. C. (2000, April). Andes: A coached learning environment for classical Newtonian physics. In *Proceedings of the 11th International Conference on College Teaching and Learning*, Jacksonville, FL. Available: www.pitt.edu/~vanlehn/icctl.pdf.

Shepard, L. A. (2000). The role of assessment in a learning culture. *Educational Researcher, 29*(7), 4–14.

Singley, M. K., & Bennett, R. E. (2002). Item generation and beyond: Applications of schema theory to mathematics assessment. In S. Irvine & P. Kyllonen (Eds.), *Item generation for test development* (pp. 361–384). Hillsdale, NJ: Lawrence Erlbaum Associates.

Tabs, E. D. (2002). *Internet access in US public schools and classrooms: 1994–2001 (NCES 2002-018)*. Washington, DC: US Department of Education, Office of Research and Improvement, National Center for Education Statistics.

Thomas, K. (2000, April 6). One school's quantum leap. *USA Today*, p. 1A. Available: www.usatoday.com/usatonline/20000406/2117463s.htm.

Tulloch, J. B. (2000). Sophisticated technology offers higher education options. *T.H.E. Journal* [On-line]. Available: www.thejournal.com/magazine/vault/A3165.cfm.

U.S. Census Bureau. (2002). Median income for 4-person families, by state. Available: www.census.gov/ftp/pub/hhes/income/4person.html.

U.S. Department of Commerce. (2002). *A nation online: How Americans are expanding their use of the Internet*. Available: http://www.ntia.doc.gov/ntiahome/dn/nationonline_020502.htm#_Toc250410.

U.S. Department of Education. (2002). *No Child Left Behind*: A desktop reference—Enhancing education through technology (II-D-1&2). Washington, DC: Author. Available: http://www.ed.gov/offices/OESE/reference/2d.html.

U.S. dominance seen slipping in Internet use, commerce. (2001). Cyberatlas: The Big Picture Geographics. Available: http://cyberatlas.internet.com/big_picture/geographics/article/0,,5911_377801,00.html.

Whalen, S. J., & Bejar, I. I. (1998). Relational databases in assessment: An application to online scoring. *Journal of Educational Computing Research, 18*, 1–13.

7

Technology, the Columbus Effect, and the Third Revolution in Learning

J. D. Fletcher
Institute for Defense Analyses

The apocryphal tale goes something like this:

> An Enlightened Leadership discovers that relatively minor investments in educational technology can significantly enhance the capabilities—productivity, competitiveness, and competence—of their domain. The next step, of course, is to appoint a "Blue-Ribbon" Committee. The task of the Committee is to design the ideal technology for education. The Committee meets, deliberates, and issues specifications for the new technology.
>
> Physically it must be rugged, lightweight, and easily portable, available anytime, anywhere. It must operate indoors and out, under a wide range of temperature, humidity, and other environmental conditions, and it must require only minimal, if any, external power support. Functionally, it must provide easy, rapid, and random access to high-quality text, black-and-white or full-color graphics, and high-resolution photographs. It must include an interface that is easily understood and usable by all, preferably communicated in natural language. It should allow self-pacing—learners should be able to proceed through instructional content as rapidly or as slowly as needed. It should be suitable for lifelong learning and readily available to a wide range of users in home, school, and workplace settings. Economically, it must be inexpensive or, as the Committee reports "requires only minimal financial investment on the part of potential end users."
>
> The Enlightened Leadership receives the Committee's report with relief. Development of the technology will require no lengthy research and development, no new taxes, no new infrastructure, and no difficult political or administrative decisions or compromises. In fact, all it will require is business as usual. The reason may be as obvious to readers here as it is to the Enlightened Leadership. The recom-

mended technology is, of course, the technology of books—already available and in place.

It seems hard to deny that writing and books effected revolutions in learning—or more precisely how we go about the business of learning. Prior to their appearance, during the earlier 100,000 or so years of human (i.e., *Homo sapiens*) existence, instruction had to take place person to person. It was expensive, slow, and produced uneven results, depending heavily on the knowledge, capabilities, and instructional expertise of the teacher. Matters concerning more than basic subsistence reached few people.

Writing developed about 7,000 years ago and progressed from picture-based ideographs to consonants and vowels represented with alphabet-based phonetization by perhaps 1000 BC. It allowed the content of advanced ideas and teaching to transcend time and place. Because of that capability, it effected a major revolution in learning. People with enough time and resources could study the words of the sages who went before them without having to rely on face-to-face interaction or the vagaries of human memory.

As discussed by Kilgour (1998), books (i.e., something beyond mud and stone tablets) were based first on papyrus and later on parchment rolls until about 300 BC, when the Romans began to sew sheets of parchment together into codices. These resembled books of today and allowed easier and random access to their content. They were also cheaper than papyrus rolls to produce because they were based on locally available parchment made from animal skin and allowed content to be placed on both sides of the sheets. Use of paper prepared from linen and cotton in about 100 AD (China) and 1200 AD (Europe) made books even less expensive. Books were by no means inexpensive, but their lowered costs made them more available to a literate and growing middle-class who, in turn, increased the demand for even more cost reductions and for greater availability of books and the learning they provided.

The full technology called for by the Committee finally became available with the introduction of books printed from moveable type (Kilgour, 1998). These printed books were first produced in China around 1000 AD and in Europe in the mid-1400s. At this point, the content of knowledge and teaching became widely and increasingly inexpensively available. The only item lacking from the Committee's list of specifications was availability of high-resolution photographs, which had to await development of photography in the mid-1800s.

The impact of writing and printed books on learning has been profound. It seems reasonable to view the emergence of writing as (among other things) the first major revolution in learning. Learning, the acquisition of knowledge, was neither inexpensive nor widely available, but it no longer required face-to-face interaction with sages. By making the content of learning, teaching, and educational material widely and inexpensively available—anytime, anywhere—the development of books printed from moveable type effected a second revolution in learning.

THE THIRD REVOLUTION IN LEARNING

At this point, we might consider an argument often advanced, but best articulated, by Clark (1983). In discussing all the means (media) we have for delivering instruction, Clark asserted that: "The best current evidence is that media are mere vehicles that deliver instruction but do not influence student achievement any more than the truck that delivers our groceries causes changes in our nutrition" (p. 445). Of course Clark had in mind more recently developed media than books, but printed books seem fair game for discussion in this context.

Are books "mere vehicles" for delivering instruction? Do they deserve more credit than that for influencing student achievement? Just as trucks are essential components in an infrastructure that has improved the nutrition of nations, so books are essential components in an infrastructure that has improved the learning, performance, and competence of people everywhere. Both trucks and books may be vehicles, but their contributions seem to be more fundamental than "mere."

However, the heart of Clark's argument remains sound. Books do not guarantee learning or student achievement. Ignorance remains plentiful, although books have appreciably diminished its supply. Clark's concern may be summed up by the notion that technology alone does not define an instructional approach—what is done with the technology matters a great deal. This point of view seems both fair and unequivocal. The presence of any technology is no guarantee that effective instructional content, effective ways to present it, or even that the unique strengths of the technology itself will be present or used. However, the absence of a technology is a reasonable guarantee that its functionalities will be missing. Without printed books, we may be back to the 1400s.

How do we improve on books? Do we have anything better? Computer technology arises as a possibility. One of the most important statements in higher order computer languages (based necessarily on what is available in every digital computer's lower order instruction set) is the "if" statement. This statement is of the (very) general form:

If <some condition> is true then do <something> otherwise do <something else>

We can well marvel at the capabilities of computers to perform millions of operations a second with perfect accuracy on the immense amounts of data that they retrieve with equal rapidity and accuracy. But what is central for this discussion is the capability of computers to adapt both the sequence and type of operations they perform based on conditions of the moment. More specifically, they can adapt the content, sequence, type, difficulty, granularity, etc. of presentations to learners and other users based on their assessment of learners' and users' current needs.

For this reason, computer technology may effect a third revolution in learning. While preserving the capabilities of writing and books to present the content of excellent instruction anytime, anywhere, they can further provide the interactions of excellent teachers, instructors, tutors, and mentors as needed by individual learners. This is not something books, movies, television, or videotape technologies can do affordably or to any appreciable degree. This interactivity is a new and significant capability. It is the core of what future commentators may view as the third revolution in learning.

DOES INTERACTIVITY MATTER?

Much of the discussion from here on centers on whether or not the instructional interactivity provided by technology matters. Can we expect a revolution in instruction equivalent to that wrought by writing and books? What can we say about the nature of this revolution and what it implies for instructional practice?

Whether interactivity matters is to some degree addressed by studies in which as much of the instruction as possible is held constant except for the level of interactivity. Two such studies were performed by Fowler (1980) and Verano (1987). Fowler compared branched presentations using computer-controlled, adaptive videodisc instruction with instruction in which the same materials were held to a fixed-content, linear sequence. She reported an effect size of 0.72 (roughly an improvement from 50th to 76th percentile performance) for ability to operate and locate faults in a movie projector, which was the objective of her instruction. Similarly, Verano compared an interactive, adaptive, branching approach for presenting instructional material with a strictly linear approach used to present identical instruction in beginning Spanish. Both of his treatments used videodisc presentations. He reported an effect size of 2.16 (roughly an improvement from 50th to 98th percentile performance) in end-of-course knowledge. These two studies, among others, suggest that interactivity—at least interactivity as defined by these studies—matters perhaps a great deal. But there is, of course, more to the story.

TUTORING AND THE INDIVIDUALIZATION OF INSTRUCTION

Individualized tutoring (one student working with one instructor) has long been viewed and used as an effective instructional procedure. Evidence of its value is found in its continued use for instruction in highly complex and high-value activities, such as aircraft piloting, advanced scientific research, and specialized medical practice. Comparisons of one-on-one tutoring with one-on-many classroom instruction might be expected to favor individualized tutoring—and they do. What is surprising about these comparisons is not the direction of their findings, but the magnitude of the differences in instructional effectiveness that they find.

Benjamin Bloom's results may be the most widely noted of these. Combining the findings of three empirical studies that compared one-on-one tutoring with classroom instruction, Bloom (1984) reported a general difference in achievement of two standard deviations (roughly an improvement from 50th to 98th percentile performance) favoring tutoring. These and similar studies suggest, on the basis of considerable empirical evidence, that differences between the results of one-on-one tutoring and classroom instruction are not just likely, but very large.

Why, then, do we not provide these manifest benefits to all our students? The answer is straightforward, obvious, and has doubtless already occurred to the reader. We cannot afford it. The issue is not effectiveness, but costs. Unless our policies toward educational funding change dramatically, we cannot afford a single tutor for every student. Bloom (1984) popularized this issue as the Two-Sigma (as in two standard deviations) Problem.

In 1975, Scriven argued that individualized instruction was an instructional imperative and an economic impossibility. Is it? Must instruction remain constrained by this reality?

Enter Moore's (famous) Law. Gordon Moore is a semiconductor pioneer and cofounder of the Intel Corporation. As recounted by Mann (2000), *Electronics* magazine interviewed Moore in 1965 and asked him about the future of the microchip industry. To make a point, Moore noted that engineers were doubling the number of electronic devices (basically transistors) on chips every year. In 1975, Moore revised his prediction to say that the doubling would occur every 2 years. If we split the difference and predict that it will occur every 18 months, our expectations fit reality quite closely. As Mann pointed out, the consequence of Moore's Law is that computers that initially sell for $3,000 will cost about half that in about 18 months.

The implication of Moore's Law for learning applications is that computers are getting exponentially less expensive and the computational capabilities we need to support instruction that is very much like individualized tutoring are becoming progressively affordable—if they are not already. The issue then becomes how we should use this increasingly affordable computational power to support learning.

We have had computer-based instruction that could tailor the content, sequence, and difficulty of instructional content to the needs of individual learners since the 1960s (e.g., Atkinson & Fletcher, 1972; Suppes, 1964), and these approaches were shown to be effective. The Stanford beginning reading programs presented on Model 33 teletypewriters running at 110 baud (about 10 characters per second) with randomly accessible digitized audio achieved effect sizes in excess of 0.80 standard deviations (Fletcher & Atkinson, 1972). Similar results were obtained for elementary school mathematics (Suppes, Fletcher, & Zanotti, 1975).

Instructional approaches used in these early programs required $2 to $3 million computers. They could easily be presented today by computers costing under $1,000.

These approaches did not seek to directly mimic the interactions that occur in human tutorial instruction. Instead they were attempts to apply results emerging from empirical studies of human cognition, memory, and learning as discussed by Suppes (1964) and Atkinson (1972). Efforts to provide tutorial dialogue emerged from approaches that were initially described as intelligent computer-assisted instruction and later as intelligent tutoring systems (Sleeman & Brown, 1982; Woolf & Regian, 2000). These have been found to be effective, occasionally yielding effect sizes in excess of 1.00 (e.g., Gott, Kane, & Lesgold, 1995). These approaches raise the question of what is it in one-on-one tutorials that accounts for their success? Can we do the same with computers?

Intensity of Instruction

This issue was discussed by Graesser and Person (1994), who compared instruction using one-on-one tutoring with classroom practice in two curriculum areas: research methods for college undergraduates and algebra for seventh graders. Tutors for the research methods course were psychology graduate students, and tutors for the algebra course were high school students. Graesser and Person found the following:

- Average number of questions teachers ask a class in a classroom hour: 3
- Average number of questions asked by any one student during a classroom hour: 0.11
- Average number of questions asked by a student and answered by a tutor during a tutorial hour:
  Research methods: 21.1
  Algebra: 32.2
- Average number of questions asked by a tutor and answered by a student during a tutorial hour:
  Research methods: 117.2
  Algebra: 146.4

Hard-core cause-and-effect is not proved by these data, but they show great differences in sheer interactivity between two approaches that also show great differences in instructional effectiveness.

Is this level of interactivity echoed by computer-based instruction? Few studies report the number of questions students using technology answer per unit of time. However, this author (Fletcher) found that K to third-grade students receiving technology-based beginning reading and arithmetic instruction on the earlier mentioned 110-baud teletypewriters were answering 8 to 12 questions a minute—questions that were individually assigned and whose answers were immediately assessed.

This level of interactivity extrapolates to 480 to 720 such questions an hour if children of this (or any) age were able to sustain this level of interaction for 60 minutes. Instead these children generally worked with the computer-based materials in daily 12-minute sessions, which extrapolates to 96 to 144 individually selected and rapidly assessed questions that these children received each day. As mentioned earlier, this computer-assisted instruction was producing effect sizes in excess of 0.80 standard deviations in comparisons with classroom instruction in both mathematics and reading.

The success of these and other computer-assisted instruction programs may have been due as much to the sheer volume of interactivity they provided as to clever instructional design or anything else. Graesser, Person, and Magliano (1995) pointed out that neither the students nor the tutors they observed were particularly sophisticated in their use of questions. Specifically, they found that the tutorial techniques long advanced by researchers and scholars—techniques such as shaping and fading (Skinner, 1968), scaffolding (Ausubel, 1960; Rogoff, 1990), reciprocal instruction (Palincsar & Brown, 1984), error diagnosis and repair (Burton, 1982; van Lehn, 1990), or advanced motivational approaches (Lepper & Woolverton, 2001)—were largely absent. About half of the questions asked by both the students and their tutors required simple yes/no responses. The techniques the tutors used were far from sophisticated, but, as the data tell us, effective. Simple approaches that aim primarily to increase interactivity may, by themselves, fill much of Bloom's two-sigma gap.

However, greater sophistication in one-on-one tutoring also pays off. Semb et al. (2000) reviewed a number of empirical studies of on-the-job training and concluded that greater knowledge and use of tutorial techniques result in greater achievement and more efficient learning. These applications are primarily found in the military and industrial world, but they are effectively one-on-one tutoring. Including advanced tutorial techniques in our computer-based tutors may allow us to exceed Bloom's two-sigma threshold. We may have just begun.

Pace of Instruction

The possibility that simple approaches may by themselves do much to fill Bloom's two-sigma gap is supported by considerations of pace—the speed with which students learn material and reach instructional objectives. Easily adjusted pacing is a capability claimed by even the most rudimentary of computer-based instruction systems.

Many teachers have been struck by the differences in the pace with which their students can learn. Consider, for instance, some findings on the time it takes for different students to reach the same instructional objectives:

- Ratio of time needed by individual kindergarten students to build words from letters: 13:1 (Suppes, 1964)
- Ratio of time needed by individual hearing-impaired and Native American students to reach mathematics objectives: 4:1 (Suppes, Fletcher, & Zanotti, 1975)
- Overall ratio of time needed by individual students to learn in Grades K to 8: 5:1 (Gettinger, 1984)
- Ratio of time needed by undergraduates in a major research university to learn features of the LISP programming language: 7:1 (A. T. Corbett, personal communication, April 30, 1998)

As with the differences between one-on-one tutoring and classroom instruction, we may not be particularly surprised to discover differences among students in the speed with which they are prepared to learn, but the magnitudes of the differences may be much larger than we expect. As we might expect from Gettinger's (1984) review, a typical K to 8 classroom will have students who are prepared to learn in 1 day what it will take other students in the same classroom 5 days to learn. This difference does not seem to be mitigated by more homogeneous grouping of students based on their abilities. The students in Corbett's (1998) university are highly selected, averaging well above 1,300 on their SATs, yet the differences in time they required to learn the fundamentals of a modestly exotic programming language remain large.

As it turns out, the differences in the speed with which different students reach given objectives may be initially due to ability, but this effect is quickly overtaken by prior knowledge as a determinant of pace (Tobias, 1989). Despite our efforts to sustain common levels of prior knowledge in classrooms by bringing every student to some minimal threshold of learning, we instead appear to increase the differences among students by about 1 year for every year they spend in elementary school (Heuston, 1997). For instance, the average spread of academic achievement in Grade 3 is about 3 years. By Grade 6, it increases to about 6 years. We are, then, working hard to make the classroom teacher's job more difficult.

The challenge this diversity presents to classroom teachers is daunting. How can they ensure that every student has enough time to reach given instructional objectives? At the same time, how can they allow students who are ready to do so surge ahead? The answer, of course, despite heroic efforts to the contrary, is that they cannot. Most classrooms contain many students who, at one end of the spectrum, are bored and, at the other end, are overwhelmed and lost.

One-on-one tutoring allows us to alleviate this difficulty by adjusting the pace of instruction to the needs and abilities of individual students. We can proceed as rapidly or as slowly as needed. We can skip what individual students have mastered and concentrate on what they have not.

As with intensity or interactivity, we do not have a direct cause-and-effect case to make for the contributions of individualized pace of instruction. But with pace as with interactivity, we find a large difference in instructional treatment associated with a large difference in instructional outcome. It does not seem unreasonable to conclude that the ability to adjust pace of instruction may also account for some of the large differences favoring individual tutoring over classroom instruction.

Again we might ask if, like one-on-one tutoring, computer-based instruction allows us to individualize the pace of instruction—pace as defined by the amount of time it takes students to reach given instructional objectives. Research findings suggest that it does. If students who could move through instructional material more quickly are prevented from doing so in classrooms, but allowed to do so in computer-based instruction, then overall we should find students reaching instructional objectives more quickly under computer-based instruction than in classrooms.

This finding arises repeatedly in reviews of instructional technology. Orlansky and String (1977) found that reductions in time to reach instructional objectives averaged about 54% across 12 evaluations of computer-based instruction used in military training. Fletcher (2002) found an average time reduction of 31% in six studies of interactive videodisc instruction applied in higher education. Kulik (1994) and his colleagues found time reductions of 34% in 17 studies of CBI used in higher education and 24% in 15 studies of adult education. These reviews are effectively independent in that they reviewed different sets of evaluation studies. From these reviews, it seems reasonable to expect reductions of about 30% in the time it takes students using computer-based instruction to reach a variety of given instructional objectives.

It is not certain that these reductions result from the speed with which students progress through fixed sets of items, from adjustments in content to take advantage of what students already know or have mastered or from some combination of these. But if we simply consider pace to be the rate with which students reach instructional objectives, then it seems reasonable to conclude that computer-based instruction reduces time to learn, as does one-on-one tutoring, primarily by not holding back students who are ready to progress.

There are three points to add to this discussion. First, the self-pacing enabled by technology-based instruction does not simply allow students to skip through content as rapidly as they can. Instead most, perhaps all, technology-based instruction includes an executive agent, which might be called an instructional management system, that allows students to progress through content as rapidly as possible, but only after they demonstrate their readiness to do so. The instructional management system ensures both instructional progress and quality in ways that books, as passive media, cannot.

Second, it turns out that 30% is a fairly conservative target. Commercial enterprises that develop technology-based instruction for the Department of Defense (DoD) regularly base their bids on the expectation that they can reduce instructional time by 50%, while holding instructional objectives constant. Noja (1991) has reported time savings through the use of technology-based instruction as high as 80% in training operators and maintenance technicians for the Italian Air Force.

Third, time saved in learning is not a trivial matter. For instance, the DoD spends about $4 billion a year on specialized skill training, which is the postbasic training needed to qualify people for the many technical jobs (e.g., wheeled vehicle mechanics, radar operators and technicians, medical technicians) needed to perform military operations. If the DoD were to reduce by 30% the time to train 20% of the people undergoing specialized skill training, it would save over $250 million per year. If it were to do so for 60% of the people undergoing specialized skill training, it would save over $700 million per year.

It is harder to assign dollar values to the time that students spend in educational settings, especially our K to 12 classrooms. This difficulty may account for the paucity of results we can find for time savings in K to 12 education. But time so spent or saved is not without cost and value. Aside from the obvious motivational issues of keeping students interested and involved in educational material, using their time well will profit both them and any society that depends on their eventual competency and achievement. The time savings offered by technology-based instruction in K–12 education could be more significant and of greater value than those obtained in posteducation training.

COST-EFFECTIVENESS

The issue for any educational decision maker faced with a unyielding budget, an unpredictable revenue stream, and unending demands for expenditures that are both urgent and imperative is not limited to instructional effectiveness. The core of such decision making is not just effectiveness, but what must be given up to get it. Most often and most specifically, this consideration centers on costs and cost-effectiveness.

Is there evidence that applications of technology in instruction are cost-effective? Despite the uncompromising need of decision makers for such evidence, little of it exists to aid their deliberations. This situation is especially prevalent among innovations, such as technology-based instruction, where researchers often seek to learn if an approach works or works better than existing practice, but very seldom to determine if it works well enough to justify its expense. The latter issue requires consideration of costs, cost models, and similar issues that researchers in instructional procedures and practices prefer to leave to others. Cost-effectiveness of an innovation is rarely considered by anyone other than the decision maker who will be pressured to adopt it.

Of course asking if an approach is cost-effective oversimplifies the issue. Cost-effectiveness is a relative term. We cannot meaningfully label some approach as

cost-effective without specifying the alternatives with which it is being compared. Cost-effectiveness studies require that a single experimental paradigm be used to compare the alternatives under consideration using comprehensive models of both costs and effectiveness. Typically, cost-effectiveness investigators either observe different levels of effectiveness achieved while holding costs constant or they observe the different costs required to reach fixed thresholds of effectiveness.

Such comparisons in technology-based instruction are hard to find—even in industrial training where all decisions are a matter of profit and loss and in military training where allocations of resources may literally be a matter of life and death. A limited cost effectiveness argument for technology-based education was reported by Fletcher (2002), who presented empirical data gathered from earlier studies by Jamison et al. (1976), Levin, Glass, and Meister (1987), and Fletcher, Hawley, and Piele (1990) to compare the costs (adjusted for inflation) of different educational interventions to raise fifth-grade mathematics scores on a standard achievement test by one standard deviation. Providing 10-minute daily sessions of computer-based instruction was found to be less expensive (and hence more cost-effective) than peer tutoring, professional tutoring, decreasing class size from 35 to either 30 or 20, or increasing the length of the school day by 30 minutes. More work of this sort is needed, but this finding suggests that a strong cost-effectiveness position for technology-based instruction is likely even at this early (relative to what may be coming) stage of development. Given that we are most likely in the horseless carriage years of the third revolution in learning, these are promising results.

INTELLIGENT TUTORING SYSTEMS

If interactivity and individualization of pace are achievable by standard approaches to computer-based instruction, is there any reason to pursue more exotic approaches? Specifically, is there any reason to develop what are called *intelligent tutoring systems* (e.g., Woolf & Regian, 2000)?

Intelligent tutoring systems may be as intelligently or unintelligently designed as any others. They involve a capability that has been developing since the late 1960s (Carbonell, 1970), but has only recently been expanding into general use. In this approach, an attempt is made to directly mimic the one-on-one dialogue that occurs in tutorial interactions. Carbonell was a computer scientist who focused on the underlying computation capabilities needed to support this approach. He contrasted ad hoc frame-oriented (AFO) approaches with information structure-oriented (ISO) approaches. Today we might be more likely to discuss "knowledge representation" as the requisite capability, but in either case the requirement is for the software to represent human knowledge—knowledge of the subject matter, knowledge of the state of the student, and knowledge of teaching strategies.

More important for those who wish to focus on instructional rather than computational capabilities are the functionalities that distinguish intelligent tutoring systems from those that have gone before. Despite current marketing efforts to describe any instructional system using technology as an intelligent instructional system, there are clear differences between what has long been the objective of these systems and what has long been available in the state of the art.

Two functionalities are critical and discriminating. First, we expect to find in an intelligent instructional system an ability to generate computer presentations and responses in real time, on demand, and as needed or requested by learners. Second, we expect to find an ability to support mixed initiative dialogue in which either the computer or the (human) student can generate, ask, and answer open-ended questions. Notably, instructional designers do not need to anticipate and prestore these interactions. The motivation for funding development of intelligent instructional systems in the early 1970s was not to apply artificial intelligence or ISO techniques to computer-based instruction, nor was it to mimic one-on-one tutorial dialogue. Instead it was to reduce the costs of instructional materials preparation by developing capabilities to generate them online in real time. Generative capabilities were intended to reduce the time and resources needed by other approaches to anticipate and prespecify all possible student–computer interactions.

Currently, intelligent tutoring systems are more sophisticated computationally and functionally than other more typical computer-based instructional systems, but they remain expensive to produce. These systems can of course adjust pace, sequence, interactivity, style, difficulty, etc. of instruction to the needs of individual learners, just as other approaches can. Notably, they can make many of the adjustments to individual learners that human tutors can. Costs to produce these systems will decrease as our techniques to develop them improve, but they may also be justified by increases in learner achievement. For instance, they show an increase in average effect size to 0.84 standard deviations (Fletcher, 2002) over the average 0.42 standard deviations (e.g., Kulik, 1994) found for other computer-based instruction approaches.

However, the main argument in favor of these systems is that they raise the bar for the ultimate effectiveness of technology-based instruction. Their unique generative and mixed initiative capabilities should eventually allow richer, more comprehensive, and more effective interactions to occur between students and the instructional system.

If Kurzweil (1999) is correct, we can expect a $1,000 unit of computing to equal the computational capability of the human brain by the year 2019 and exceed it thereafter. Computers may then become more effective in providing instruction than human tutors even if humans use all the techniques Graesser et al. (1995) found they now neglect. We may not be implanting integrated circuits in our brains as Kurzweil suggested we might by 2029. However, using computers to discover more than any human agent can about the unique potential of every individual, and then devising effective and individualized procedures to reach it, seems both an appealing and realistic prospect.

Whatever the case, the extensive tailoring of instruction to the needs of individual students that can be obtained through the use of generative, intelligent tutoring systems can only be expected to increase. Our current approaches may be reaching their limits. Intelligent tutoring systems may make available far greater instructional effectiveness and efficiencies than we can obtain from the approaches we are using now.

THE COLUMBUS EFFECT IN INSTRUCTIONAL TECHNOLOGY

Prognostications aside, these technological approaches to instruction may provide yet another example of what might be called the Columbus Effect. As readers will recall, Columbus sailed west intending to find India (and a lucrative spice route). Instead he (re)discovered what became a new world for Europeans. Such a result typifies technological progress. Seeking one thing based on familiar, common practice, we inevitably end up with something else, unforeseen and unexpected. Wireless telegraph produced something functionally quite different than the telegraph—namely, radio. Similarly, efforts to make a carriage run without a horse produced automobiles—to say nothing of gas stations, motels, and the Santa Monica Freeway. Seeking affordable one-on-one tutoring through automation, we may end up with something no one now envisions. The metaphor based on current practice gets us started. The result may surprise us all.

As we begin with a vision of one-on-one tutoring made affordable through the use of computing and telecommunications technologies, our work may center on efforts to mimic the interactions that occur between human tutors and their students. We may be pursuing humanless tutoring—just as books, television, and other noninteractive media may be viewed humanless lecturing, but we are likely to end up with something quite different—in function, appearance, and use.

Some hint of what this different result may be is perhaps seen in the vision promulgated by the advanced distributed learning (ADL) initiative currently pursued jointly by the Department of Defense and the White House Office of Science and Technology Policy. This vision is based on the expectation that most, if not all, human knowledge will become available as shareable, interoperable objects in the World Wide Web. The ADL initiative seeks specification and development of these objects, but it does so because it envisions something that might be called a personal learning associate (PLA).

Physically, a PLA will be a computer that is either carried or worn. In keeping with the suggestions of Kurzweil among others, it may even be implanted to provide direct brain–computer interaction, although that possibility seems more remote and subject to more review and consideration than the current discussion requires.

Functionally, a PLA will provide wireless connection to the Web or its successor in the global communication ether. A student of any age will use it for learning or a decision maker, such as an electronics technician, military tactician, or business planner, will use it to help solve practical problems. All will interact with it using spoken natural language. The PLA will provide a full range of display capabilities including text, graphics, and photographs as specified by the apocryphal Committee's specifications. It will also provide animation, digital video of some sort augmented by a full range of high-fidelity sound, and perhaps tactile and haptic feedback as well. Olfaction remains under review.

Most of the physical capabilities of the PLA are in the state of the art. Many of the software capabilities are also available or, like shareable courseware objects, soon will be (Fletcher & Dodds, 2000). Its instructional and decision-aiding functionalities remain longer term and more elusive. These functionalities call for the PLA to develop and then use comprehensive, intimate, and accurate knowledge of the student/user to identify, collect, and integrate shareable instructional objects. This process will be accomplished on demand, in real time, and be precisely tailored to the individual's needs, capabilities, interests, and cognitive style. If the intention is to help solve a problem, information the PLA provides will not just be expert, but delivered in a form that the individual is prepared to understand. If the intention is to establish a more permanent change in the individual's cognitive ability (i.e., to bring about learning), it will do so efficiently and effectively in ways far superior to those we now imagine.

Obviously, we have a way to go to realize this vision. What form it eventually takes, how it is used, what infrastructure it engenders, and what impact it has on our lives all remain to be seen, but its key capabilities may well arise from the intelligent instructional systems we are now learning how to build. The Columbus Effect will kick in sooner or later, but beginning with a guiding metaphor based on individualized, tutorial instruction and mentoring seems as good a way as any to advance. The goal of learning to do something that is within our reach but outside our grasp has long been a stimulus for human progress.

IMPACT ON RESEARCH AND THEORY

Reviews by Krendl and Lieberman (1988) and Schacter and Fagnano (1999) continue to echo earlier recommendations by Suppes (1964) and others to apply advances in cognitive and learning theory to the development of technology-based instruction. Such efforts will improve the quality of instruction delivered. More important, they will provide feedback to theories of cognition and learning about where they are right, where they might use some improvement, and where they have left serious gaps that need to be filled. This is the traditional interplay of theory and empirical research that has served other areas of systematic investigation so well. Technology-based instruction, with its precise control over stimulus inputs and equally precise measurement of response outputs, should play a unique role in completing the feedback loop between theory and empirical research.

Such feedback will produce significant advances not just in instruction, but in related areas as well. Instructional applications using technology test the notions—or theories—of human cognition, learning, and instruction embodied in them in at least two ways. First, the ability to put a proposed notion, model, or theory into a computer algorithm is a significant demonstration by itself. If a notion cannot be captured by an algorithm, it may not be testable; if it is not testable, it is not worth serious consideration.

Second, an instructional application is an instantiation that tests the correctness of a notion, model, or theory. To the extent that an application achieves its goals, the notion(s), model(s), or theory of cognition, learning, and instruction on which it is based may be viewed as correct. However, tests of instructional applications seldom yield simple answers. The fine-grain data that technology-based instruction makes available will increase the richness of feedback we receive. Few theories will be shown to be perfectly and thoroughly correct. The more complete diagnostic information concerning where they are correct, where they are not, and what they lack is critical. Detailed and specific information of this sort produces significant advances in other fields, and we should expect nothing less from our instructional applications.

THE ENGINEERING OF INSTRUCTION

Beyond issues of feedback for theory, our third revolution in learning may effect a shift in instruction from art to engineering. We tend to view teaching as an intensely human activity—something that at best can only be accomplished by human teachers interacting with human students. Master teachers do exist, and many of us have benefited from the attentions of at least one teacher who, if not a master, was at least able to impress on us the value, benefits, and pleasure of scholarship. But such occasions appear to be more the exception than the rule. Our perceptions of teaching as a warmly experienced, human activity are at variance with much that we experience in real classrooms, where each student is one of many waiting for that portion of instruction that addresses his or her individual needs.

To a great extent, successful instruction is a matter of design—the creation of an environment to maximize the probability that learning will occur and that specified instructional objectives will be achieved by every student. Mostly we seek and test for accuracy of knowledge. However, we may also seek other objectives such as speed of response, retention, insight, and transfer of knowledge as well as continued interest in and respect for the subject matter. These goals may be compatible to some extent, but at some point they diverge and require different approaches that compete with each other for classroom resources of human energy, funding, and, especially, time.

The design of such an environment may be viewed as an art—a highly personal, hit-or-miss affair. To be fair to our students and productive in our instruction, we need to establish a science of design that allows these many instructional objectives to be accomplished by many different hands. In short, we need an engineering of instruction in which specific designs reliably yield specific instructional outcomes.

The notion of instruction as engineering may well be unpopular. Fox (1994) noted that, "One of the more difficult problems in dealing with improvement in public education is to replace the notion of teaching as an art form with that of instructional delivery as a systems science" (p. 2). In classroom instruction, sufficiently precise control of instruction environments may be out of the question. Tailoring learning environments to the needs, interests, and capabilities of individual students can be achieved, at substantial expense, by one-on-one tutoring. As suggested earlier, technology can make such engineering of learning environments affordable. However, it may also do it better.

This possibility is suggested by Meehl's studies of decision making. These studies were intended to determine the degree to which human judgment would be needed to augment purely algorithmic (linear regression) predictions of such outcomes as patient responses to treatment and graduate school applicants' academic success. Instead, as described by Dawes (1971), the statistical prediction "floor" turned out to be a ceiling. In all 20 cases reviewed by Meehl, statistical predictions based on straightforward (algorithmic) linear models turned out to be superior to the judgments of humans, even though the linear models were derived solely from human decisions. The same superiority may obtain for decision making by technology-based tutors.

IMPACT ON OUR CURRICULAR GOALS

In 1960, T. F. Gilbert wrote:

> If you don't have a gadget called a 'teaching machine,' don't get one. Don't buy one; don't borrow one; don't steal one. If you have such a gadget, get rid of it. Don't give it away, for someone else might use it. This is the most practical rule, based on empirical facts from considerable observation. *If you begin with a device of any kind, you will try to develop the teaching program to fit that device.* (p. 478; italics original)

Gilbert may be both right and wrong. He is certainly correct in suggesting that instructional designers and developers who adapt a teaching machine will try to fit

the teaching program to it. The new functionalities such a device makes available motivate its adaptation in the first place. One can imagine students long ago poring over clay tablets or papyrus rolls once their teachers learned how to design teaching programs to take advantage of written language. The same may be said for printed books, which, simply by being more accessible and less expensive than papyrus rolls or codices, allowed teachers different assumptions in the development of their teaching programs. The same is doubtlessly true for our third instructional revolution, which applies interactive, computer-based technology to the problems and processes of instruction.

It is less certain that such adaptations, fitting the teaching program to the device, is a significant evil to be avoided at all costs. Certainly if technology causes us to remove or de-emphasize essential elements of our teaching programs, it will diminish their effectiveness. It is also true that if properly applied, technology will improve, if not revolutionize, the effectiveness and efficiency of our teaching programs. It is up to researchers, developers, and instructors, not the technology itself, to see that it does.

In either case, the application of technology will change both what and how we teach. Technology will raise the bar for our curricular aspirations. Tutorial simulations will afford our students experiences, access to exotic (i.e., expensive and unavailable) devices, and immersion in collaborative problem solving that we could not provide in any other way. Intelligent tutoring capabilities will permit tutorial interactions or simple conversations with experts and expertise that would otherwise be out of the question. Because these interactions will be geared to each student's level of ability and prior knowledge, they will produce levels of understanding that would otherwise be unattainable. Sooner or later we will be forced by necessity, public pressure, or our own professional integrity to adapt our teaching programs to the new functionalities technology makes available. The most important consequence of the third revolution in instruction may not be that it improves efficiency and effectiveness of what we do now, but that it will change what it is that we choose to do.

SUMMARY

There is more to be discussed about the third revolution in learning wrought by technology, but that must await wiser, better informed commentary. A few statements can be made about this revolution with modest certainty:

- It will make the functionalities of individualized tutoring widely accessible and affordable.
- It will permit interactive, individualized learning to take place anytime, anywhere.

- It will (eventually) bring about profound changes in our educational institutions and the roles and responsibilities of people (teachers, students, and administrators) in them.
- It will help bring into being a nation of lifelong learners who are prepared to meet the challenges of the new century and thrive in the global marketplace.
- Thanks to the Columbus effect, it will lead to capabilities, uses, and functionalities of which we are now only dimly, if at all, aware.
- It will produce radical change in the practice and processes of instruction.

Basically there is just one way that people learn, most likely involving growth or chemical changes in the synapses. The third revolution in learning will not change learning at this level any more than writing or books did. But substantially increasing the probability that such fundamental changes will occur across all manner of humans in all manner of environments, does seem to qualify as a radical change—one that warrants the term *revolution*. Because it will increase the tempo with which learning occurs, it might well be called a revolution in learning—the magnitude of which we have only seen twice before in human history.

REFERENCES

Atkinson, R. C. (1972). Ingredients for a theory of instruction. *American Psychologist, 27*, 921–931.

Atkinson, R. C., & Fletcher, J. D. (1972). Teaching children to read with a computer. *The Reading Teacher, 25*, 319–327.

Ausubel, D. P. (1960). The use of advanced organizers in the learning and retention of meaningful verbal material. *Journal of Educational Psychology, 51*, 267–272.

Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*, 4–16.

Burton, R. R. (1982). Diagnosing bugs in a simple procedural skill. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 157–183). New York: Academic Press.

Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems, 11*, 190–202.

Clark, R. E. (1983). Reconsidering research on learning from media. *Review of Educational Research, 53*, 445–459.

Dawes, R. M. (1971). A case study of graduate admissions: Application of three principles of human decision-making. *American Psychologist, 26*, 180–188.

Fletcher, J. D. (2002). Evidence for learning from technology-assisted instruction. In H. F. O'Neil, Jr. & R. Perez (Eds.), *Technology applications in education: A learning view* (pp. 79–99). Hillsdale, NJ: Lawrence Erlbaum Associates.

Fletcher, J. D., & Atkinson, R. C. (1972). An evaluation of the Stanford CAI program in initial reading (Grades K through 3). *Journal of Educational Psychology, 63*, 597–602.

Fletcher, J. D., & Dodds, P. V. W. (2000). All about ADL. *Learning Circuits*. An Internet publication of the American Society for Training & Development: http://www.learningcircuits.org/may2000/.

Fletcher, J. D., Hawley, D. E., & Piele, P. K. (1990). Costs, effects, and utility of microcomputer assisted instruction in the classroom. *American Educational Research Journal, 27*, 783–806.

Fowler, B. T. (1980). *The effectiveness of computer-controlled videodisc-based training*. Unpublished doctoral dissertation, University of Iowa, Iowa City, IA (University Microfilms No. 8114254).

Fox, R. G. (1994). *A systems approach to education*. Warrenton, VA: Learning Technology Institute.

Gettinger, M. (1984). Individual differences in time needed for learning: A review of the literature. *Educational Psychologist, 19*, 15–29.

Gilbert, T. F. (1960). On the relevance of laboratory investigation of learning to self-instructional programming. In A. A. Lumsdaine & R. Glaser (Eds.), *Teaching machines and programmed learning: A source book* (pp. 475–485). Washington, DC: National Education Association of the United States

Gott, S. P., Kane, R. S., & Lesgold, A. (1995). *Tutoring for transfer of technical competence* (AL/HR-TP-1995-0002). Brooks AFB, TX: Armstrong Laboratory, Human Resources Directorate.

Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal, 31*, 104–137.

Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology, 9*, 1–28.

Heuston, D. H. (1997). *School improvement models: The manual model and the speed of light*. Sandy, UT: Waterford Institute.

Jamison, D. T., Fletcher, J. D., Suppes, P., & Atkinson, R. C. (1976). Cost and performance of computer-assisted instruction for education of disadvantaged children. In J. T. Froomkin, D. T. Jamison, & R. Radner (Eds.), *Education as an industry* (pp. 201–240). Cambridge, MA: Ballinger.

Kilgour, F. G. (1998). *The evolution of the book*. New York: Oxford University Press.

Krendl, K. A., & Lieberman, D. A. (1988). Computers and learning: A review of recent research. *Journal of Educational Computing Research, 4*, 367–389.

Kulik, J. A. (1994). Meta-analytic studies of findings on computer-based instruction. In E. L. Baker & H. F. O'Neil, Jr. (Eds.), *Technology assessment in education and training* (pp. 9–33). Hillsdale, NJ: Lawrence Erlbaum Associates.

Kurzweil, R. (1999). *The age of spiritual machines: When computers exceed human intelligence*. New York: Viking Press.

Lepper, M. R., & Woolverton, M. (2001). The wisdom of practice: Lessons learned from the study of highly effective tutors. In J. Aronson (Ed.), *Improving academic achievement: Contributions of social psychology* (pp. 133–156). Orlando, FL: Academic Press.

Levin, H. M., Glass, G. V., & Meister, G. R. (1987). Cost-effectiveness of computer-assisted instruction. *Evaluation Review, 11*, 50–71.

Mann, C. C. (2000). The end of Moore's Law? *Technology Review*. Available: http://www.techreview.com/articles/may00/mann.htm.

Noja, G. P. (1991). DVI and system integration: A further step in ICAI/IMS technology. In R. J. Seidel & P. R. Chatelier (Eds.), *Advanced technologies applied to training design* (pp. 161–189). New York: Plenum.

Orlansky, J., & String, J. (1977). *Cost effectiveness of computer-based instruction in military training* (IDA Paper P-1375). Arlington, VA: Institute for Defense Analyses.

Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction, 1*, 117–175.

Rogoff, B. (1990). *Apprenticeship in thinking: Cognitive development in social contexts*. New York: Oxford University Press.

Schacter, J., & Fagnano, C. (1999). Does computer technology improve student learning and achievement? How, when, and under what conditions. *Journal of Educational Computing Research, 20*, 329–343.

Scriven, M. (1975). Problems and prospects for individualization. In H. Talmage (Ed.), *Systems of individualized education* (pp. 199–210). Berkeley, CA: McCutchan.

Semb, G. B., Ellis, J. A., Fitch, M., & Kuti, M. B. (2000). On-the-job training (OJT): Theory, research, and practice. In S. Tobias & J. D. Fletcher (Eds.), *Training & retraining: A handbook for business, industry, government, and the military* (pp. 289–311). New York: Macmillan Reference.

Skinner, B. F. (1968). *The technology of teaching*. New York: Appleton, Century, Crofts.

Sleeman, D., & Brown, J. S. (1982). *Intelligent tutoring systems*. New York: Academic Press.

Suppes, P. (1964). Modern learning theory and the elementary-school curriculum. *American Educational Research Journal, 1*, 79–93.

Suppes, P., Fletcher, J. D., & Zanotti, M. (1975). Performance models of American Indian students on computer-assisted instruction in elementary mathematics. *Instructional Science, 4*, 303–313.

Tobias, S. (1989). Another look at research on the adaptation of instruction to student characteristics. *Educational Psychologist, 24*, 213–227.

van Lehn, K. (1990). *Mind bugs: The origins of procedural misconceptions*. Cambridge, MA: MIT Press.

Verano, M. (1987). *Achievement and retention of Spanish presented via videodisc in linear, segmented, and interactive modes*. Unpublished doctoral dissertation, University of Texas, Austin, TX (DTIC No. ADA 185 893).

Woolf, B. P., & Regian, J. W. (2000). Knowledge-based training systems and the engineering of instruction. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 339–356). New York: Macmillan Reference.

III

Affordances of Software